


Metrics Correlation and Analysis Service

A workflow-driven system for retrieval, correlation, analysis, and display of multisource heterogeneous operational data.

Ted Hesselroth , Gabriele Garzoglio, Tanya Levshina, Parag Mhashilkar, Andrew Baranovski
{tdh, garzoglio, tlevshin, parag}@fnal.gov
Computing Division, Fermilab, Batavia, IL



Monitoring
Dashboard
Data Mining

Aggregation
Correlation
Analysis

Sources from:

HTML
XML
RRD
database
text
images

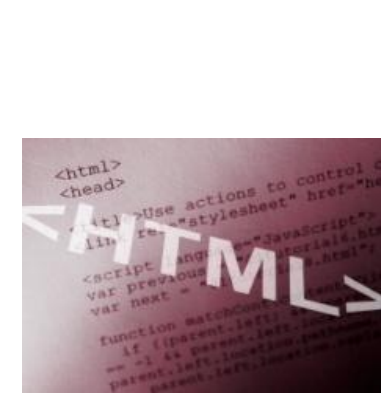
MCAS

Collect Store Display


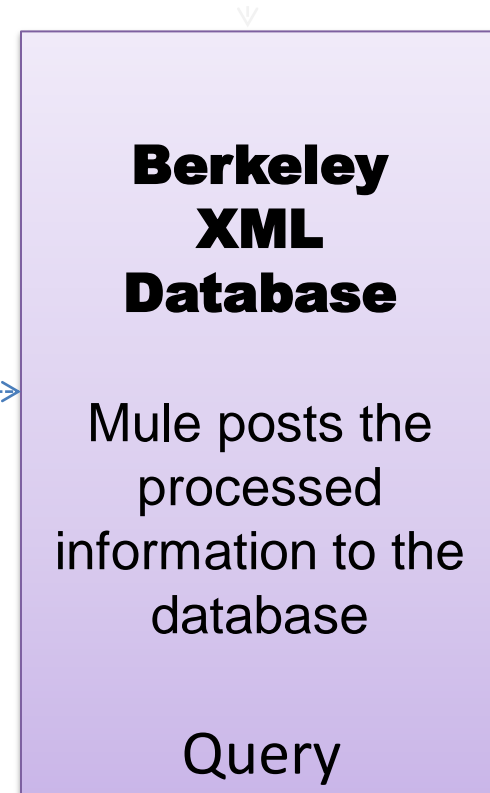
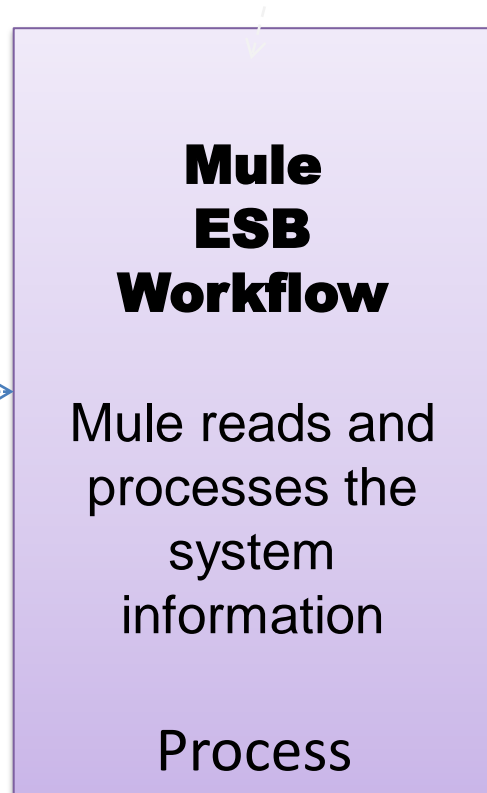
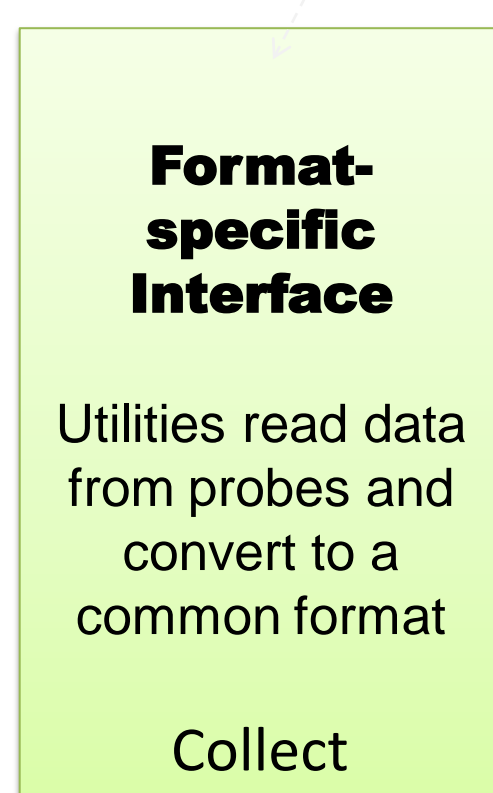
Process Query

MCAS Components

Operational Data



Operational data is collected by probes, administrative interfaces, log scraping, etc.



User Display



Operator or system administrator views unified, coherent, customizable display of system state and history.

Triggered Workflows

Data can be either pushed or pulled into the XML database of MCAS. In the push model, probes periodically post XML containing system information to a service endpoint provided by the Mule instance. Mule then processes the data in a workflow, for example, converting the schema or aggregating data from multiple sources. In the pull model, probe data may be acquired through periodic queries performed by Mule, or may be triggered from user actions in the graphical interface.

Workflow Configuration

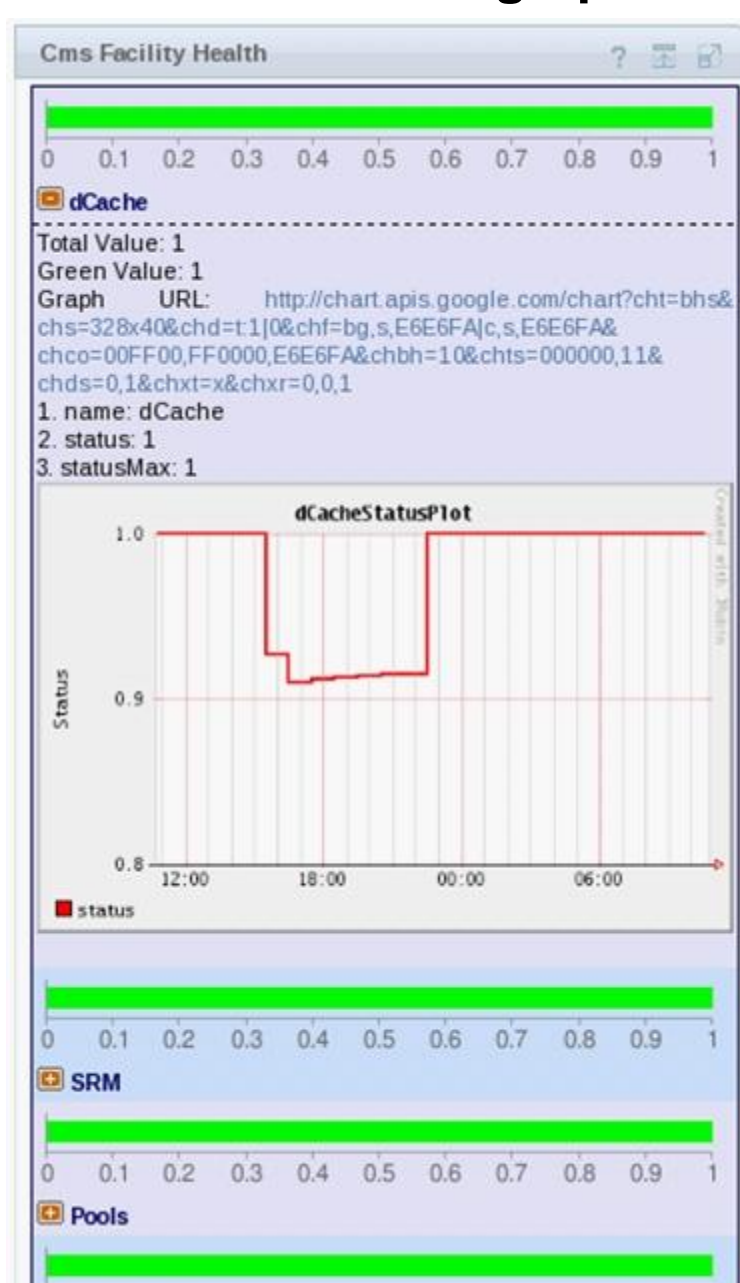
The transformations that occur in workflows are based on models that are specific to each operational system. MCAS provides tools to manage data source and workflow characteristics. In particular, the characteristics of the workflow, including its dependencies, are specified as parameters in python objects. The workflow configuration tool maps these objects into runnable Mule workflow configurations.

Database Maintenance

Automatic rotation of database files is built into the software. When a database reaches a specified size or number of entries, a new instance of it is created. Previous data from rotated versions is still transparently accessible. This allows gradual deletion or archiving of the oldest, non-used data without affecting the functionality of the software. Automatic backup is also supported.

Use Cases

dCache Health Bargraph

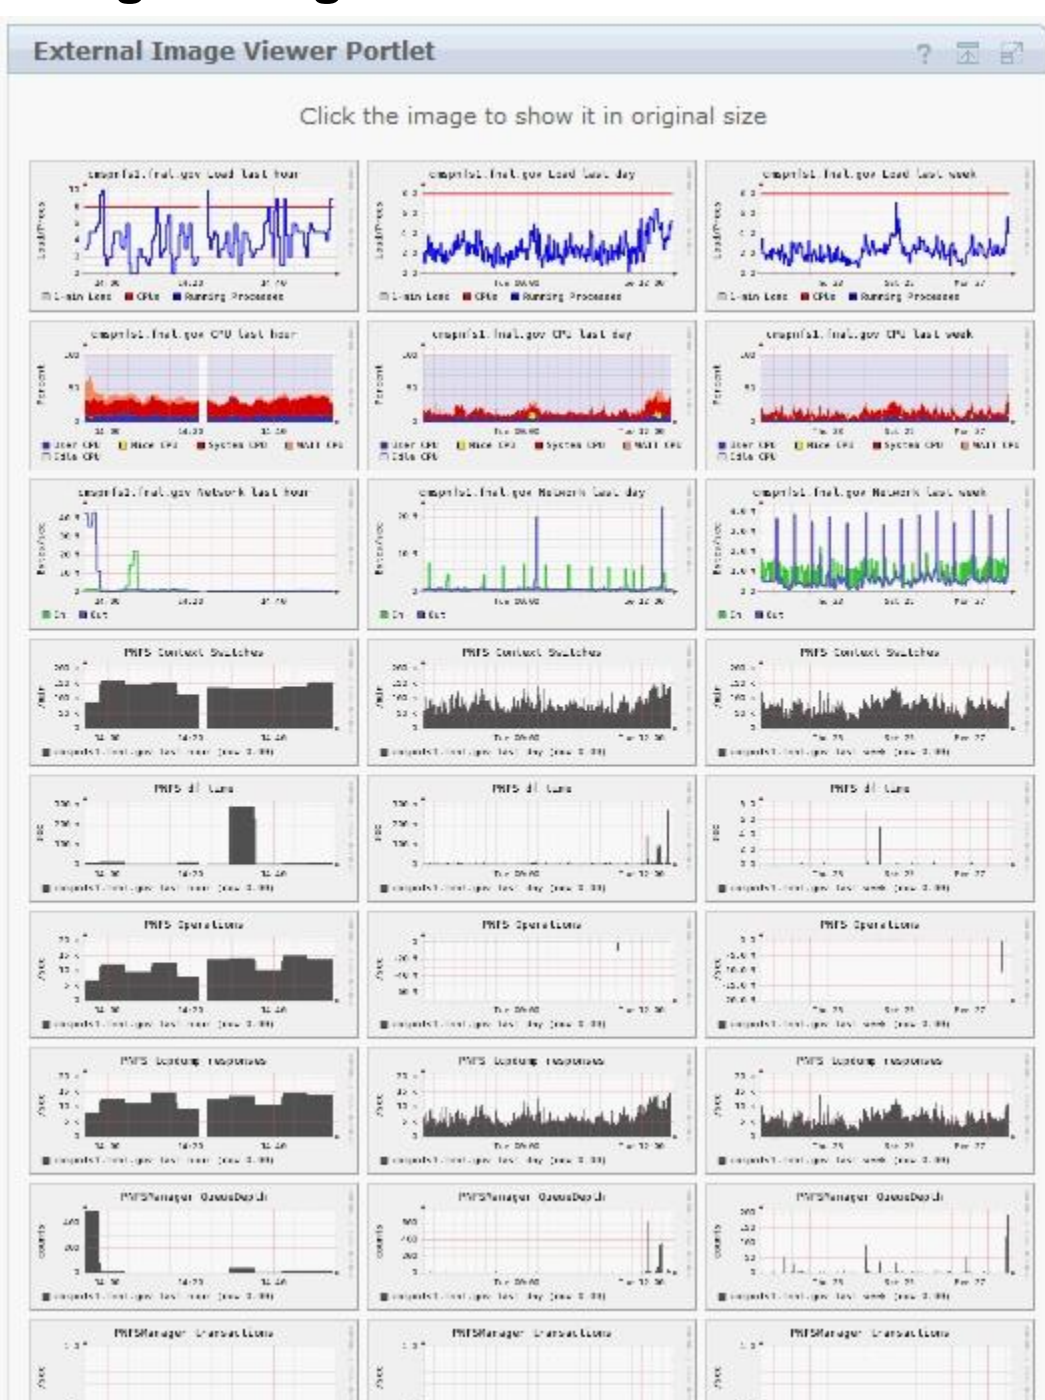


Cell Status Table

Cell	Domain	up/down	ping time	version	offline
cd-cache-262-11	v-cmsor262-110Domain	up	1.0	1.0	false
cd-cache-262-12	v-cmsor262-120Domain	up	1.0	1.0	true
cd-cache-262-13	v-cmsor262-130Domain	up	1.0	1.0	true
cd-cache-262-14	v-cmsor262-140Domain	up	1.0	1.0	true
cd-cache-262-15	v-cmsor262-150Domain	up	1.0	1.0	true
cd-cache-262-16	v-cmsor262-160Domain	up	1.0	1.0	true
cd-cache-262-17	v-cmsor262-170Domain	up	1.0	1.0	true
cd-cache-262-18	v-cmsor262-180Domain	up	1.0	1.0	true
cd-cache-262-19	v-cmsor262-190Domain	up	1.0	1.0	true
cd-cache-262-20	v-cmsor262-200Domain	up	1.0	1.0	true
cd-cache-262-21	v-cmsor262-210Domain	up	1.0	1.0	true
cd-cache-262-22	v-cmsor262-220Domain	up	1.0	1.0	true
cd-cache-262-23	v-cmsor262-230Domain	up	1.0	1.0	true
cd-cache-262-24	v-cmsor262-240Domain	up	1.0	1.0	true
cd-cache-262-25	v-cmsor262-250Domain	up	1.0	1.0	true
cd-cache-262-26	v-cmsor262-260Domain	up	1.0	1.0	true
cd-cache-262-27	v-cmsor262-270Domain	up	1.0	1.0	true
cd-cache-262-28	v-cmsor262-280Domain	up	1.0	1.0	true
cd-cache-262-29	v-cmsor262-290Domain	up	1.0	1.0	true
cd-cache-262-30	v-cmsor262-300Domain	up	1.0	1.0	true
cd-cache-262-31	v-cmsor262-310Domain	up	1.0	1.0	true
cd-cache-262-32	v-cmsor262-320Domain	up	1.0	1.0	true
cd-cache-262-33	v-cmsor262-330Domain	up	1.0	1.0	true
cd-cache-262-34	v-cmsor262-340Domain	up	1.0	1.0	true
cd-cache-262-35	v-cmsor262-350Domain	up	1.0	1.0	true
cd-cache-262-36	v-cmsor262-360Domain	up	1.0	1.0	true
cd-cache-262-37	v-cmsor262-370Domain	up	1.0	1.0	true
cd-cache-262-38	v-cmsor262-380Domain	up	1.0	1.0	true
cd-cache-262-39	v-cmsor262-390Domain	up	1.0	1.0	true
cd-cache-262-40	v-cmsor262-400Domain	up	1.0	1.0	true
cd-cache-262-41	v-cmsor262-410Domain	up	1.0	1.0	true
cd-cache-262-42	v-cmsor262-420Domain	up	1.0	1.0	true
cd-cache-262-43	v-cmsor262-430Domain	up	1.0	1.0	true
cd-cache-262-44	v-cmsor262-440Domain	up	1.0	1.0	true
cd-cache-262-45	v-cmsor262-450Domain	up	1.0	1.0	true
cd-cache-262-46	v-cmsor262-460Domain	up	1.0	1.0	true
cd-cache-262-47	v-cmsor262-470Domain	up	1.0	1.0	true
cd-cache-262-48	v-cmsor262-480Domain	up	1.0	1.0	true
cd-cache-262-49	v-cmsor262-490Domain	up	1.0	1.0	true
cd-cache-262-50	v-cmsor262-500Domain	up	1.0	1.0	true
cd-cache-262-51	v-cmsor262-510Domain	up	1.0	1.0	true
cd-cache-262-52	v-cmsor262-520Domain	up	1.0	1.0	true
cd-cache-262-53	v-cmsor262-530Domain	up	1.0	1.0	true
cd-cache-262-54	v-cmsor262-540Domain	up	1.0	1.0	true
cd-cache-262-55	v-cmsor262-550Domain	up	1.0	1.0	true
cd-cache-262-56	v-cmsor262-560Domain	up	1.0	1.0	true
cd-cache-262-57	v-cmsor262-570Domain	up	1.0	1.0	true
cd-cache-262-58	v-cmsor262-580Domain	up	1.0	1.0	true
cd-cache-262-59	v-cmsor262-590Domain	up	1.0	1.0	true
cd-cache-262-60	v-cmsor262-600Domain	up	1.0	1.0	true
cd-cache-262-61	v-cmsor262-610Domain	up	1.0	1.0	true
cd-cache-262-62	v-cmsor262-620Domain	up	1.0	1.0	true
cd-cache-262-63	v-cmsor262-630Domain	up	1.0	1.0	true
cd-cache-262-64	v-cmsor262-640Domain	up	1.0	1.0	true
cd-cache-262-65	v-cmsor262-650Domain	up	1.0	1.0	true
cd-cache-262-66	v-cmsor262-660Domain	up	1.0	1.0	true
cd-cache-262-67	v-cmsor262-670Domain	up	1.0	1.0	true
cd-cache-262-68	v-cmsor262-680Domain	up	1.0	1.0	true
cd-cache-262-69	v-cmsor262-690Domain	up	1.0	1.0	true
cd-cache-262-70	v-cmsor262-700Domain	up	1.0	1.0	true
cd-cache-262-71	v-cmsor262-710Domain	up	1.0	1.0	true
cd-cache-262-72	v-cmsor262-720Domain	up	1.0	1.0	true
cd-cache-262-73	v-cmsor262-730Domain	up	1.0	1.0	true
cd-cache-262-74	v-cmsor262-740Domain	up	1.0	1.0	true
cd-cache-262-75	v-cmsor262-750Domain	up	1.0	1.0	true
cd-cache-262-76	v-cmsor262-760Domain	up	1.0	1.0	true
cd-cache-262-77	v-cmsor262-770Domain	up	1.0	1.0	true
cd-cache-262-78	v-cmsor262-780Domain	up	1.0	1.0	true
cd-cache-262-79	v-cmsor262-790Domain	up	1.0	1.0	true
cd-cache-262-80	v-cmsor262-800Domain	up	1.0	1.0	true
cd-cache-262-81	v-cmsor262-810Domain	up	1.0	1.0	true
cd-cache-262-82	v-cmsor262-820Domain	up	1.0	1.0	true
cd-cache-262-83	v-cmsor262-830Domain	up	1.0	1.0	true
cd-cache-262-84	v-cmsor262-840Domain	up	1.0	1.0	true
cd-cache-262-85	v-cmsor262-850Domain	up	1.0	1.0	true
cd-cache-262-86	v-cmsor262-860Domain	up	1.0	1.0	true
cd-cache-262-87	v-cmsor262-870Domain	up	1.0	1.0	true
cd-cache-262-88	v-cmsor262-880Domain	up	1.0	1.0	true
cd-cache-262-89	v-cmsor262-890Domain	up	1.0	1.0	true
cd-cache-262-90	v-cmsor262-900Domain	up	1.0	1.0	true
cd-cache-262-91	v-cmsor262-910Domain	up	1.0	1.0	true
cd-cache-262-92	v-cmsor262-920Domain	up	1.0	1.0	true
cd-cache-262-93	v-cmsor262-930Domain	up	1.0	1.0	true
cd-cache-262-94	v-cmsor262-940Domain	up	1.0	1.0	true
cd-cache-262-95	v-cmsor262-950Domain	up	1.0	1.0	true
cd-cache-262-96	v-cmsor262-960Domain	up	1.0	1.0	true
cd-cache-262-97	v-cmsor262-970Domain	up	1.0	1.0	true
cd-cache-262-98	v-cmsor262-980Domain	up	1.0	1.0	true
cd-cache-262-99	v-cmsor262-990Domain	up	1.0	1.0	true
cd-cache-262-100	v-cmsor262-1000Domain	up	1.0	1.0	true

Storage administrators want to see the status of each dCache cell, displayed in a sortable table. For example, this table allows the sorting of information in such a way that all offline cells appear on top. More sorting criteria can be added.

Ganglia Images




The administrators of the Minos experiment collect multiple graphs of various kinds (e.g. number of running/idle jobs per user, read data rates on bluearc over time, etc). The MCAS image viewer portlet allows to organizing and displaying multiple graphs on one page as thumbnails. A end-user can enlarge a graph by clicking on it.

Door Status Heatmap



The administrators of public dCache running at Fermilab include in the dashboard the health status of gridftp doors over time. MCAS can display this information as a heatmap, where the color of each bin represents the status of a particular door during some period of time, e.g., the percentage of the time the door was up or down.

RRD Integration



MCAS has support for storing information in a Round Robin Database (RRD). Internally, MCAS can use JRobin as RRD. When used in the push model, probes periodically post Data in RRD format to the RRD service endpoint provided by Mule. Mule then processes the data in a workflow to store the data in one or more RRDs. In the pull model, probe data can be retrieved in RRD or XML format using appropriate queries. Users can create complex workflows that store the data in RRD databases and pull the data to create plots that are displayed using the Image Viewer portlet.